

Document: T10/01-227r1
To: T10 Committee Membership
From: Edward A. Gardner, Ophidian Designs
Subject: SRP Buffer Descriptor Formats

Date: August 3, 2001

Revision 1: changes from July 19-20 SRP working group. Format and size moved to bytes 5-7, bit mask of requested / supported formats exchanged during login.

At present SRP defines two buffer descriptor formats, direct and indirect. Both are specifically designed to work with Infiniband™ and VI, and perform data transfers on the same RDMA channel used to transfer commands and status.

However, if a proposal such as 01-085r1 were adopted, we would need additional information to perform a data transfer, specifically some identification of the RDMA data channel(s) to use for data transfer. There is ongoing discussion whether a scatter/gather data transfer should be obliged to use a single RDMA channel group for all segments of the data transfer or might be allowed to use a separate RDMA channel group for each segment. I believe there are sound arguments in favor of both modes of operation.

Another issue, discussed at the June 19-20 SRP working group, is that the current buffer descriptor formats are designed to work with Infiniband™ and VI. We do not know the requirements of other transport protocols that might be used in the future. They might require differently formatted buffer descriptors or multiple formats of buffer descriptor.

This proposal is that we restructure the current fields that define buffer descriptor formats to allow more than two formats. Our current two formats (direct and indirect) will be defined, all others will be reserved.

A future multi-channel data transfer proposal would define one or more additional buffer descriptor formats. One benefit is that we do not need to add reserved fields to SRP_CMD as proposed by 01-085r2, since any additional fields can be defined as part of the new buffer descriptor format. In particular this eliminates any risk that the size of any reserved fields we add now might turn out to be wrong for the final proposal.

Similarly, if some future transport might need a differently sized buffer descriptor, we could address it at the time the issue arises by defining a different buffer descriptor format.

Proposal specifics.

At present buffer descriptors are defined by a one bit indirect flag and an eight bit count field. The indirect flag indicates whether the buffer descriptor is direct or indirect. The count field indicates the number of 16-byte memory descriptors present in the buffer descriptor.

These would be replaced by a buffer format code and a size field. I propose that the buffer format code be four bits and the size field be eight bits. The minimum size is probably two bits for the format code and four bits for the size field. Combining the two fields into a single field (e.g. eight bits) is also practical.

The size field will indicate the length of the buffer descriptor in 4-byte units. Permissible combinations of the format code and size fields would be:

Format=0, size=0: no buffer descriptor present.

Format=0, size=4: direct buffer descriptor present.

Format=1, size=n*4+5: indirect buffer descriptor present containing a PARTIAL MEMORY DESCRIPTOR LIST of n memory descriptors.

All other combinations would be reserved.

The fields that determine buffer descriptor format are currently in SRP_CMD bytes 1, 2 and 3. These will be moved to bytes 5, 6 and 7. The rationale is that bytes 1, 2 and 3 are reserved and available in all information units. We would like to keep those available for future control fields that might be useful for other information units, not just SRP_CMD.

SRP_CMD byte 1 currently contains DOIND in bit 7 and DIIND in bit 6. These will be changed to DATA-OUT BUFFER DESCRIPTOR FORMAT in byte 5 bits 7-4 and DATA-IN BUFFER DESCRIPTOR FORMAT in byte 5 bits 3-0.

SRP_CMD byte 2 currently contains DOCOUNT. It will be changed to DATA-OUT BUFFER DESCRIPTOR SIZE in byte 6.

SRP_CMD byte 3 currently contains DICOUNT. It will be changed to DATA-IN BUFFER DESCRIPTOR SIZE in byte 7.

The INDREQ bit in SRP_LOGIN_REQ and the INDSUP bit in SRP_LOGIN_RSP will be renamed and expanded to sixteen bit fields. REQUIRED BUFFER FORMATS in bytes 24-25 of SRP_LOGIN_REQ is a bit mask of the buffer formats that may be specified by the initiator. The target shall reject the login if it is unable to support all of the requested formats. SUPPORTED BUFFER FORMATS in bytes 24-25 of SRP_LOGIN_RSP and SRP_LOGIN_REJ is a bit mask of the buffer formats supported by the target.

The detailed wording for the above changes appears in T10/01-230r1.

Note: I do not recall whether the working group explicitly discussed returning SUPPORTED BUFFER FORMATS for a rejected login as well as an accepted login. To me it seems the proper thing to do. Anyone who objects should raise this issue on the reflector or at the next working group meeting.